

09/966,263  
6000/446

98-C17-TAP

## INTELLIGENT DATA STORAGE MANAGER

### Field of the Invention

This invention relates to data storage subsystems and, in particular, to a dynamically mapped virtual data storage subsystem which includes a data storage manager that functions to combine the non-homogeneous physical devices contained in the data storage subsystem to create a logical device with new and unique quality of service characteristics that satisfy the criteria for the policies appropriate for the present data object.

### Problem

It is a problem in the field of data storage subsystems to store the ever increasing volume of application data in an efficient manner, especially in view of the rapid changes in data storage characteristics of the data storage elements that are used to implement the data storage subsystem and the increasingly specific need of the applications that generate the data.

Data storage subsystems traditionally comprised homogeneous collections of data storage elements on which the application data was stored for a plurality of host processors. As the data storage technology changed and a multitude of different types of data storage elements became available, the data storage subsystem changed to comprise subsets of homogeneous collections of data storage elements, so that the application data could be stored on the most appropriate one of the plurality of subsets of data storage elements. Data storage management systems were developed to route the application data to a selected subset of data storage elements and a significant amount of processing was devoted to ascertaining the proper data storage destination for a particular data set in terms of the data storage characteristics of the selected subset of data storage elements. Some systems also migrate data through a hierarchy of data storage elements to account for the timewise variation in the data storage needs of the data sets.

In these data storage subsystems, the quality of service characteristics are determined by the unmodified physical attributes of the data storage elements that are used to populate the data storage subsystem. One exception to this rule is disclosed in U.S. Patent No. 5,430,855 titled "Disk Drive Array Memory System Using Nonuniform Disk Drives," which discloses a data storage subsystem that uses an array of data

storage elements that vary in their data storage characteristics and/or data storage capacity. The data storage manager in this data storage subsystem automatically compensates for any nonuniformity among the disk drives by selecting a set of physical characteristics that define a common data storage element format. However, the data storage utilization of the redundancy groups formed by the data storage manager is less than optimal, since the least common denominator data storage characteristics of the set of disk drives is used as the common disk format. Thus, disk drive whose data storage capacity far exceeds the smallest capacity disk drive in the redundancy group suffers from loss of utilization of its excess data storage capacity. Therefore, most data storage subsystems do not utilize this concept and simply configure multiple redundancy groups, with each redundancy group comprising a homogeneous set of disk drives. A problem with such an approach is that the data storage capacity of the data storage subsystem must increase by the addition of an entire redundancy group. Furthermore, the replacement of a failed disk drive requires the use of a disk drive that matches the characteristics of the remaining disk drives in the redundancy group, unless loss of the excess data storage capacity of the newly added disk drive were incurred, as noted above.

Thus, it is a prevalent problem in data storage subsystems that the introduction of new technology is costly and typically must occur in fairly large increments, occasioned by the need for the data storage subsystem to be comprised of homogeneous subset of data storage devices, even in a virtual data storage subsystem. Therefore, data administrators find it difficult to cost effectively manage the increasing volume of data that is being generated in order to meet the needs of the end users' business. In addition, the rate of technological innovation is accelerating, especially in the area of increases in data storage capacity and the task of incrementally integrating these new solutions into existing data storage subsystems is difficult to achieve.

#### Solution

The above described problems are solved and a technical advance achieved by the present intelligent data storage manager that functions to combine the non-homogeneous physical devices contained in a data storage subsystem to create a logical device with new and unique quality of service characteristics that satisfy the criteria for the policies appropriate for the present data object. In particular, if there

is presently no logical device that is appropriate for use in storing the present data object, the intelligent data storage manager defines a new logical device using existing physical and/or logical device definitions as component building blocks to provide the appropriate characteristics to satisfy the policy requirements. The intelligent data storage manager uses weighted values that are assigned to each of the presently defined logical devices to produce a best fit solution to the requested policies in an n-dimensional best fit matching algorithm. The resulting logical device definition is then implemented by dynamically interconnecting the logical devices that were used as the components of the newly defined logical device to store the data object.

#### Brief Description of the Drawing

Figure 1 illustrates in block diagram form the overall architecture of a data storage subsystem in which the present intelligent data storage manager is implemented;

Figure 2 illustrates a three-dimensional chart of the operating environment of the present intelligent data storage manager;

Figure 3 illustrates one example of a virtual device that can be configured by the present intelligent data storage manager; and

Figure 4 illustrates a three-dimensional chart of a user policy that must resolve priorities between two attributes: Cost per MB, and Time to First Byte.

#### Detailed Description

Figure 1 illustrates in block diagram form the overall architecture of a data storage subsystem 100 in which the present intelligent data storage manager 110 is implemented. The data storage subsystem is connected to a plurality of host processors 111-114 by means of a number of standard data channels 121-124. The data channels 121-124 are terminated in a host interface 101 which provides a layer of name servers 131-134 to present virtual implementations of existing defined physical device interfaces to the host processors 111-114. As far as the host processors 111-114 are concerned, the name servers 131-134 implement a real physical device. The name servers 131-134 convert the user data received from the host processor 111-114 into a user data object which can be either converted into a canonical format or left in binary format. The object handle server maps the object handle to logical device addresses and allows multiple instances of a data object. The object handle server 102 maps the user data object into a data

space for storage. The mapping is determined by the policies programmed into the policy manager 105 of the data storage subsystem 100 and subject to security layer 103. The persistent storage for the object space is determined by the logical device manager 104 which allocates or creates a logical device based upon policies for storing the user data object. A logical device is a composite device and can consist of a real physical device such as a tape 151, a disk 152, optical disk 153, another logical device, such as Logical Device 1 which comprises a RAID 5 disk array 154, Logical Device N which comprises middleware software 155 that accesses another logical device, such as access of a logical device over a network connection, or combinations of the above. The logical device definition abstracts the nature of the real device associated with the persistent storage. The changes implemented in the technology of the persistent storage are thereby rendered transparent to the host application.

If there is presently no logical device that satisfies the criteria for the policies appropriate for a user data object, the logical device manager 104 creates a new logical device definition with the appropriate data storage characteristics to satisfy the policy requirements using existing physical and/or logical device definitions as component building blocks. The logical device manager 104 uses weighted values that are assigned to each of the presently defined logical devices to produce a best fit solution to the requested policies in an n-dimensional best fit matching algorithm. Thus, the intelligent data storage manager 110 maps the virtual device to the user data object rather than mapping a data object to a predefined data storage device. The various data storage attributes that are used by the intelligent data storage manager 110 to evaluate the appropriateness of a particular virtual device include, but are not limited to: speed of access to first byte, level of reliability, cost of storage, probability of recall, and expected data transfer rate. The logical device manager 104 stores the mapping data which comprises a real time definition of the available storage space in the data storage subsystem 100. Once one of the current logical device definitions meet the criteria required by a data object, the logical device manager 104 either allocates space on an existing instance of a logical device of that type or creates a new instance of that type of logical device.

#### Policy Attributes

The policy attributes and the potential algorithms that are used to map user requirements to storage devices are managed by the intelligent storage manager

110. A typical general set of attributes for storage devices is shown in Table 1:

**Table 1: Policy Attributes**

Name of Attribute	Range of Values (Dimension)
Cost per MB (lg)	\$0.0001 to \$1000.00
Time to first byte (lg)	Ns to days
Random read	0.0001 to 1000 MB/sec
Random write	0.0001 to 1000 MB/sec
Sequential read	0.0001 to 1000 MB/sec
Sequential write	0.0001 to 1000 MB/sec
Sequential (tape) or random (disk) storage or recall	0 to 10 (where: 0= sequential, 10= random)
Size (lg)	Bytes to petabytes
Probability of recall	0 to 10 (where: 0= lowest, 10= highest)
Virtual or real Device	yes/no
Level of reliability	0 to 10 (where: 0= minimum, 10= 100%)
Others to be defined...	

Each of these attributes has a range or dimension of "values". Each dimension needs to be relatively uniform in its number scheme. For example, each dimension could have a numeric value for 0.0 to 10.0. Some dimensions need to be logarithmic (lg) because of the inherent nature of the dimension. For example, Cost per MB can be defined as a logarithmic dimension that runs from the \$0.001 for tape storage to \$10s for RAM. So one approach is to do a distance calculation of the difference between the customer's policy requirements and each storage device's policy attributes. In addition, levels of priority among attributes can be specified since certain dimensions may be more important than others (reliability, for example). When the intelligent storage manager 110 must resolve between

conflicting priority levels, the logical storage manager 104 tries to find ways to combine single devices into an optimal, logical device using logical combining operators.

#### Operation of the Intelligent Data Storage Manager

The present intelligent data storage manager 110 is responsive to one of the host processors 111 initiating a data write operation by transmitting a predefined set of commands over a selected one of the communication links to the data storage subsystem 100. These commands include a definition of the desired device on which the present data object is to be stored, typically in terms of a set of data storage characteristics. Figure 2 illustrates a three-dimensional (of the above-noted multiple dimensions) chart of the operating environment of the present intelligent data storage manager 110 and the location of the host specified data storage device with respect to this environment. In particular, as mapped in a Cartesian coordinate system, the cost, data transfer rate, and data access time comprise the three axes used to measure the performance characteristics of the various physical 151-153 and virtual 154-155 devices of the data storage subsystem 100. As shown in Figure 3, the standard tape 151, disk 152, and optical 153 devices each have a set of defined characteristics that can be mapped to the three-dimensional space of Figure 2. The user has requested that their data be stored on a device, whose data storage characteristics do not match the data storage characteristics of any of the devices presently defined in the data storage subsystem 100. The desired data storage characteristics are shown mapped as a locus in the three-dimensional space in Figure 2. The intelligent data storage manager 110 must therefore map the existing set of physical devices that are contained in the data storage subsystem 100 to satisfy the desired set of data storage characteristics defined by the user. This problem comprises a three-dimensional best fit mapping process wherein the set of available physical and virtual devices are mapped to match or at least approximate the desired set of data storage characteristics. This is accomplished by creating a composite virtual device that implements the defined desired data storage characteristics. For example, assume that the user has requested a data storage device that has a 20MB/sec read performance and the data storage subsystem 100 is equipped with 5MB/sec tape drives as one of the types of physical devices. The intelligent data storage manager 110 can create a 20MB/sec data storage device by configuring a Redundant Array of Inexpensive

Tape drives (RAIT) to connect a plurality of the existing tape drives 151 in parallel to thereby achieve the desired data throughput.

#### Examples of Operation of the Intelligent Data Storage Manager

There are many instances of data file storage where the needs of the application and/or user do not correspond to the reality of the data storage characteristics of the various data storage elements 151-153 and virtual data storage elements 154-155 that are available in the data storage subsystem 100. For example, the application "video on demand" may require a high reliability data storage element and fast access to the initial portion of the file, yet not require fast access for the entirety of the file since the data is typically read out at a fairly slow data access rate. However, the required data transfer bandwidth may be large, since the amount of data to be processed is significant and having a slow speed access device as well as a narrow bandwidth would result in unacceptable performance. Furthermore, the cost of data storage is a concern due to the volume of data. The intelligent data storage manager 110 must therefore factor all of these data storage characteristics to determine a best fit data storage device or devices to serve these needs. In this example, the defined data storage characteristics may be partially satisfied by a Redundant Array of Inexpensive Tapes since the reliability of this data storage device is high as is the data bandwidth, yet the cost of implementation is relatively low, especially if the configuration is a RAIT-5 and the data access speed is moderate. In making a determination of the appropriate data storage device, the intelligent data storage manager 110 must review the criticality of the various data storage characteristics and the amount of variability acceptable for that data storage characteristic.

#### Defining Attribute Values

All devices support some form of quality of service, which can be described as attributes with certain fixed values. For example, they cost \$xxx per megabyte of data or have nnn access speed. The intelligent storage manager 110 provides an algorithmic way to use these attributes to determine the perfect device, as specified by user policy. In some cases, the perfect device is a logical device that is constructed when the intelligent storage manager 110 rank orders the distance between 1) how the user would like to have data stored and 2) the storage devices that are available. This logical device can span both disk and tape subsystems and, therefore, blurs the distinction between disk and tape.

The diagram of Figure 4 shows an example of a user policy that must resolve priorities between two attributes: Cost per MB, and Time to First Byte. To resolve this, the intelligent storage manager 110 could create a logical device that is the mixture of disk and tape that best conforms to the specific policies the user has requested. In this example, some data could be stored on disk for quick access and some data could be stored on tape for lower cost of storage. Or the intelligent storage manager 110 could create a policy that migrates a small file between disk and tape over time: after a week the file would be transferred to tape to lower storage cost.

Table 2 provides a more complex comparison of device attributes versus attributes defined through user policy. In this example, the set of attributes of the following storage subsystems: single disk, RAID, single tape drive, and RAIT are listed. The intelligent storage manager 110 determines an optimal storage solution by doing a distance calculation between 1) the set of attributes for each device and 2) the set of attributes for a file (defined through user policy).

For example, the calculation below denotes the vector for point P by  $[x_1(P), x_2(P), x_3(P)]$ . Then the distance between points 1 and 2 is

$$\sqrt{[(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2]}$$

Where:       $x_1$  is the attribute value defined by user policy.  
                 $x_2$  is the attribute value defined for the device.

This example is for three dimensions. To extend it to more dimensions, take the difference between corresponding components of the two vectors, square this difference, add this square to all the other squares, and take the square root of the sum of the squares. Of course, you don't need to do the square root if you're simply looking for the point closest to a give point.

Table 2: Device attributes vs. user-defined policy

Device	Cost/MB	Time to first byte	MB/sec read	MB/sec write	Sequential or Random	Reliability
Disk	0.15	12ms	3MB/sec	3MB/sec	5	1
RAID	10.00	6ms	80MB/sec	20MB/sec	3	3
Tape	.001	30sec	5MB/sec	5MB/sec	0	2
RAIT	.005	40sec	20MB/sec	20MB/sec	0	4

**User-defined policy (per attribute)**

File	.01	1 sec or less	.1 MB/sec or less	.1 MB/sec or less	0	3
------	-----	---------------	-------------------	-------------------	---	---

In the present example, the realized data storage device can be a composite device or a collection of composite devices. For example, the video on demand file data storage requirements can be met by the virtual device illustrated in Figure 3. The virtual device 300 can comprise several elements 301, 302, each of which itself comprises a collection of physical and/or virtual devices. The virtual device 300 comprises a first device 301 which comprises a set of parallel connected disk drives 310-314 that provides a portion of the data storage capability of the virtual device 300. These parallel connected disk drives 310-314 provide a fast access time for the application to retrieve the first segment of the video on demand data to thereby provide the user with a fast response time to the file request. The bulk of the video on demand data file is stored on a second element 302 that comprises a Redundant Array of Inexpensive Tapes device that implements a RAIT-5 storage configuration. The relative data storage capacity of the two data storage elements 301, 302 is determined by the amount of data that must be provided to the user on

a priority basis and the length of time before the remainder of the file can be staged for provision to the user.

#### Time Analysis

The data storage manager 110 implements devices that support some form of quality of service. These attributes have some type of fixed value: they cost so much- they have XX access speed. The data storage manager 110 can also rank order the distances between how the user wishes to have a data file stored compared to the storage devices that are in the data storage subsystem 100. From this the data storage manager 110 can also come up with some alternative storage methods- for example, the data storage manager 110 can do a mixture of disk and tape to achieve the qualities that the user is looking for. The data storage manager 110 can put some of the data file on disk for quick access and some of it on tape for cheap storage as noted above. Another alternative factor is if there is a file that the user wants stored at a certain \$\$ per megabyte, it can be migrated from disk to tape over a certain period of weeks and the average cost of storage complies with the user policy definition. So, the data storage manager 110 must evaluate quickly what devices are available and the data storage manager 110 compares how the user wants to store the data file. If the data storage manager 110 doesn't have a perfect match, the mixtures of devices are rank ordered and investigated to try and achieve the policy that is defined by the user.

#### Summary

The intelligent data storage manager functions to combine the non-homogeneous physical devices contained in a data storage subsystem to create a logical device with new and unique quality of service characteristics that satisfy the criteria for the policies appropriate for the present data object. The intelligent data storage manager uses weighted values that are assigned to each of the presently defined logical devices to produce a best fit solution to the requested policies in an n-dimensional best fit matching algorithm. The resulting logical device definition is then implemented by dynamically interconnecting the logical devices that were used as the components of the newly defined logical device to store the data object.

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

**BLACK BORDERS**

**IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**

**FADED TEXT OR DRAWING**

**BLURRED OR ILLEGIBLE TEXT OR DRAWING**

**SKEWED/SLANTED IMAGES**

**COLOR OR BLACK AND WHITE PHOTOGRAPHS**

**GRAY SCALE DOCUMENTS**

**LINES OR MARKS ON ORIGINAL DOCUMENT**

**REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**

**OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**